

**UNITED STATES DISTRICT COURT
WESTERN DISTRICT OF TEXAS
AUSTIN DIVISION**

Intellectual Ventures I LLC and Intellectual Ventures II LLC, Plaintiffs v. VMware, Inc., Defendant.	Civil Action No. 1:19-cv-01075-ADA
---	------------------------------------

PLAINTIFFS' RESPONSIVE CLAIM CONSTRUCTION BRIEF

TABLE OF CONTENTS

Introduction.....	1
The '686 and '726 Patents	2
1. Disputed Terms in the '686 Patent.....	4
A. Modif[y/ied] [a] resource allocation / modify[ing] [the] computer resources allocated to a virtual server ('686 patent claims 5-7)	4
B. “resource unavailable messages” / “denied requests to modify a resource allocation” ('686 patent claims 5-7)	6
C. “determination that a virtual server is overloaded” ('686 patent claims 5-7).....	10
D. “virtual server” ('686 patent claims 5-7)	12
E. “determining that a second physical host can accommodate the requested modified resource allocation” ('686 patent claims 5-7)	14
2. Disputed Terms in the '726 Patent.....	16
A. “resource denials” ('726 patent claims 1, 4-5 & 8)	16
B. “quality of service guarantee” ('726 patent claims 1 & 4)	17
3. Alleged Means Plus Function Terms for the '726 and '686 Patents	18
A. '686 patent claim 7 “component” terms (i.e., clauses 1-3 of Ex. A)	18
B. '726 patent claim 1, 4 & 5 (i.e., clauses 4-8 of Ex. A)	21
The '752 Patent.....	25
4. Disputed Terms in the '752 Patent.....	25
A. “exhausted” ('752 patent claims 1, 9 and 24)	25
B. “consumed” ('752 patent claims 1, 9 and 24).....	28
C. “service” ('752 patent claims 1, 3, 9 and 24).....	29
D. Means-Plus-Function Terms.....	29
The '051 Patent	31
5. Disputed Terms in the '051 Patent.....	31
A. “virtual server” (claims 1, 3, and 6).....	31
B. “physical interface[s]” (claims 1 and 3).....	33
C. “storing . . .” / “storing . . .” / “receiving . . .” / “determining . . .” / “determining . . .” / “sending . . .” / “using . . .” (claims 1 and 3)	34
D. “customer forwarding [table(s)/information]” (claims 1 and 3)	37
6. Disputed Terms in the '818 Patent.....	38
A. “hierarchical token bucket resource allocation” / “token(s)” (claims 1, 17, 30, 32 and 42)	38

B. “enforcing . . .” / “receiv[e/ing] . . .” / “classify[ing] . . .” / “compar[e/ing] . . .” / “forward[ing] . . .” / “buffer[ing] . . .” (claims 1, 17, 30, 32, 33, 37-39, 42).....	40
C. “maintain[ing] a connection over a network fabric” (claims 1, 17, 30, 32, 42)	42
D. “virtual [network/storage network] interface layer of an application server” (claims 1, 17, 32, 42)	44
E. Alleged Means Plus Function Elements (claim 17) (individually set forth in Ex. C).....	46
Conclusion	48

TABLE OF AUTHORITIES

Cases

<i>Accent Packaging, Inc. v. Leggett & Platt, Inc.</i> , 707 F.3d 1318 (Fed. Cir. 2013).....	36, 45
<i>ActiveVideo Networks Inc. v. Verizon Communications Inc.</i> , 694 F.3d 1312 (Fed. Cir. 2012).....	1
<i>ACTV, Inc. v. Walt Disney Co.</i> , 346 F.3d 1082 (Fed. Cir. 2003).....	35
<i>Arlington Indus., Inc. v. Bridgeport Fittings, Inc.</i> , 632 F.3d 1246 (Fed. Cir. 2011).....	36, 41, 43
<i>Baldwin Graphic Sys., Inc. v. Siebert, Inc.</i> , 512 F.3d 1338 (Fed. Cir. 2008).....	35
<i>Bell & Howell Document Mgmt. Prods. Co. v. Altek Sys.</i> , 132 F.3d 701 (Fed. Cir. 1997).....	38
<i>Comark Commc'ns, Inc. v. Harris Corp.</i> , 156 F.3d 1182 (Fed. Cir. 1998).....	11, 14, 17, 29
<i>Cox Commc'ns, Inc. v. Sprint Commc'n Co. LP</i> , 838 F.3d 1224 (Fed. Cir. 2016).....	42
<i>Duncan Parking Techs, Inc. v. IPS Grp., Inc.</i> , 914 F.3d 1347 (Fed. Cir. 2019).....	17
<i>Fisher-Rosemount Sys. v. ABB Ltd.</i> , Case No. 4:18-cv-00178, 2019 WL 6830806 (S.D. Tex. Dec. 12, 2019)	18, 46
<i>Georgetown Rail Equip. Co. v. Holland L.P.</i> , 867 F.3d 1229 (Fed. Cir. 2017).....	41
<i>Graphon Corp. v. Autotrader.com, Inc.</i> , Case No. 2:05-cv-530 (TJW), 2007 WL 1870622 (E.D. Tex. June 28, 2007)	15
<i>iFly Holdings LLC v. Indoor Skydiving Germany GmbH</i> , Case No. 2:14-cv-1080-JRG-RSP, 2015 WL 9258264 (E.D. Tex. Dec. 18, 2015)	13
<i>Inline Plastics Corp. v. EasyPak, LLC</i> , 799 F.3d 1364 (Fed. Cir. 2015).....	6
<i>Innova/Pure Water, Inc. v. Safari Water Filtration Sys., Inc.</i> , 381 F.3d 1111 (Fed. Cir. 2004).....	40
<i>Liebel-Flarsheim Co. v. Medrad, Inc.</i> , 358 F.3d 898 (Fed. Cir. 2004).....	10
<i>Medrad, Inc. v. MRI Devices Corp.</i> , 401 F.3d 1313 (Fed. Cir. 2005).....	10, 33
<i>Medtronic, Inc. v. Guidant Corp.</i> , 465 F.3d 1360 (Fed. Cir. 2006).....	35

<i>O2 Micro Int'l Ltd. v. Beyond Innovation Tech. Co.,</i> 521 F.3d 1351 (Fed. Cir. 2008).....	1, 42
<i>Paragon Sols., LLC v. Timex Corp.,</i> 566 F.3d 1075 (Fed. Cir. 2009).....	41
<i>Renishaw PLC v. Marposs Societa' per Azioni,</i> 158 F.3d 1243 (Fed. Cir. 1998).....	33, 46
<i>Sound View Innovations, LLC v. Facebook, Inc.,</i> Case No. 16-cv-116 (RGA), 2017 WL 2221177 (D. Del. May 19, 2017).....	22
<i>Telcordia Tech., Inc. v. Cisco Sys., Inc.,</i> 612 F.3d 1365 (Fed. Cir. 2010).....	19, 47
<i>Thorner v. Sony Computer Entm't Am. LLC,</i> 669 F.3d 1362 (Fed. Cir. 2012).....	10, 33, 43
<i>VLSI Tech. LLC, v. Intel Corp.,</i> Case No. 1:19-CV-00977-ADA.....	30
<i>Zeroclick, LLC v. Apple Inc.,</i> 891 F.3d 1003 (Fed. Cir. 2018).....	46
Other Authorities	
MPEP § 2173.05(e) (8 th ed. Rev. 2 May 2004).....	15

Introduction

VMware reveals its claim construction strategy with the first line of its Opening Claim Construction Brief (“VMware Br.”): “This case is immensely complicated.” *VMware Br.* at 1. In fact, VMware has made the claim construction process unnecessarily complicated in an effort to create non-infringement arguments. VMware reads limitations into the asserted claims and conjures up invalidity positions with a kitchen-sink approach to claim construction. VMware employs a number of tactics to further these efforts. VMware’s proposed constructions (1) mix and match terms within claims and even across patents; (2) seek to construe entire claim elements, as well as multiple elements, as single proposed “terms;” (3) argue entire groups of terms should be construed as means-plus-function despite the terms not including the word “means;” and (4) utilize a fifty-three (53) page expert declaration to skew claim language. Below, and in its Opening Claim Construction Brief (“IV Br.”), IV demonstrates that by following well-established claim construction canons, each of the claims and their constituent terms can be well understood by examining the intrinsic record.

VMware misinterprets Federal Circuit law to repeatedly argue that the Court should adopt VMware’s constructions for terms that should simply be construed as plain and ordinary meaning. The case relied on by VMware, *O2 Micro*, stands for the proposition that a term should not be construed with its ordinary meaning if there are multiple possible ordinary meanings or if plain and ordinary will not resolve the parties’ dispute. *O2 Micro Int’l Ltd. v. Beyond Innovation Tech. Co.*, 521 F.3d 1351, 1361 (Fed. Cir. 2008). That is not the case here. IV’s plain and ordinary proposals resolve the claim construction dispute with VMware and are consistent with the intrinsic record and how a person of ordinary skill in the art (POSITA) would have understood them at the time of the invention. *See ActiveVideo Networks Inc. v. Verizon Commc’ns Inc.*, 694 F.3d 1312, 1325-26 (Fed. Cir. 2012) (Federal Circuit rejects argument under *O2 Micro* that a term should not have been construed with its plain and ordinary meaning because doing so did, in fact, resolve the claim construction dispute between the parties).

The patents-in-suit utilize language readily understood by one of ordinary skill in the art. Examining each term individually in the context of the claim in which it appears bears this out.

The '686 and '726 Patents

VMware's efforts to overly complicate this case begin with its approach to United States Patent Nos. RE 44,686 ("the '686 patent) and RE 42,726 ("the '726 patent"). For example, despite the term "quality of service guarantee" appearing in only select claims of the '726 patent, VMware attempts to read it into the construction of nearly all claims across both the '686 and '726 patents. In another example, it points to use of the term "virtual server" in United States Patent No. RE 43,051 ("the '051 patent") to support its construction of the same term in the '686 and '726 patents, but then argues that a different construction of "virtual server" is required in the '051 patent itself.

To help clarify the following arguments with respect to the '686 and '726 patents a short background is instructive. The '686 patent and '726 patent are related and share a specification. They have the same inventors, and both originated from Ensim Corporation. The '726 patent is a reissue of U.S. Pat. No. 6,985,937 and the '686 patent is a continuation of the application that would issue as the '726 patent. Despite sharing a specification, there are key differences in the claims that should be considered. For instance, the '726 patent includes claims with the term "quality of service guarantee" whereas the '686 patent does not include that term at all. VMware glosses over this important distinction. Another difference is that the '726 patent's claims are written in a way that focuses more on components, e.g., 'a virtual server resource monitor,' 'a virtual server resource modifier,' 'a load balancing module,' and a 'dynamic virtual server mover,' and how those components operate together. The '686 patent's claims on the other hand, are written to focus more on the messages and indications that are used by the components to come to the desired result. In some cases, this leads to different terms having similar meanings, like 'resource denials' from the '726 patent and 'denied requests to modify a resource allocation' from the '686 patent. Such a result is not inconsistent, however, because it is based on the specific language and sequence of the claims at issue.

A short background of the '051 patent-in-suit is also helpful in exposing the needless complication VMware creates by cross-referencing the '051 patent in arguments regarding the '686 and '726 patents. While all three patents originated with Ensim Corporation, they are unrelated and share no common inventors. In addition, the inventions taught by the '686 and '726 are quite different than that taught by the '051 patent.

The core concept embodied in the '686 and '726 patents concerns the dynamic modification of resource allocations made to a virtual server by moving the virtual server from one physical host (a server) to another physical host if the virtual server is not getting enough resources to function properly on its current host. Through a series of components, indications and messages, the inventors of the '726 and '686 patents conceived of a way to allow the physical resources used by a virtual server to be dynamically changed if its physical host was overloaded and could not allocate additional resources to that virtual server. That way is to dynamically transfer that virtual server to another host. Because re-distributing such resources on a dynamic basis requires coordination and cooperation among many different components, the sequence of the claimed elements is of great importance in these patents.

The '051 patent, while also dealing with virtualized servers, utilizes virtualization to a different end. The '051 patent's specific combination of virtualized components, tunnels, and physical interfaces allow data centers to be siloed on a per-tenant basis so that each tenant may use the same physical hardware as another tenant without risk of security breach or hardware failure. Furthermore, because the tenant's servers are virtualized, the network addresses of one tenant's virtual servers may overlap with the network addresses of another tenants' virtual server, which further lessens the risk of a security breach while providing each tenant with a lot of IP addresses.

It is against these distinct backdrops that the claims of these three patents must be analyzed. When examined in the proper context, it becomes clear that IV's proposed constructions are appropriate, and that VMware is attempting to needlessly sow confusion.

1. Disputed Terms in the '686 Patent

A. **Modif[y/ied] [a] resource allocation / modify[ing] [the] computer resources allocated to a virtual server ('686 patent claims 5-7)**

IV's Proposed Construction	VMware's Proposed Construction
“modif[y/ied] set of functions and features of a physical host(s) used in implementing tasks for each virtual server” / “modify[ing] a set of the functions and features of a physical host(s) used in implementing tasks for each virtual server”	“modif[y/ied] [a] quality of service guarantee” / “modify[ing] [the] quality of service guarantee of a virtual server”

In its Opening Brief, VMware mischaracterizes the intrinsic record and IV's proposed construction in an attempt to read unsupported limitations into the disputed terms. First, VMware mischaracterizes the prosecution history of the '686 patent by stating that during prosecution of the '937 patent (parent of the '686 patent), the applicant disclaimed any meaning of “resource allocation” other than “quality of service guarantee.” *See* VMware Br. at 3. Second, VMware tries to support this “quality of service guarantee” construction with support from one of the claims of the '726 patent, despite the fact that this phrase is not found in any of the claims of the '686 patent. Third, in arguing that IV's construction reads out “allocation” and seeks to construe only “modifying a resource” VMware is in effect disregarding the wording of the claim. *See id.*

VMware's argument that there was an “unmistakable prosecution history disclaimer” requires a misreading of the '686 patent's prosecution history. The applicant's statement, relied on by VMware, that “a resource allocation for a process is specified as a quality of service guarantee,” referred to the specific claim language of the embodiment captured by then claim 1, not to the scope of the invention as a whole (and certainly not to claims entirely devoid of the term “quality of service guarantee”).

Claim 1, as written at the time of the applicants' statement, is reproduced below:

1. In a computer system including one or more physical hosts, a method for dynamically modifying the computer resources allocated to a process, the method comprising:

determining that a resource allocated to a process on a first physical host is overloaded; and

in response to the resource being overloaded, increasing a quality of service guarantee for the process.

(emphasis added). As can be seen by viewing the prosecution history, the claim then being prosecuted required that the resource allocation be specified as the claimed quality of service guarantee.

A closer look at the evolution of the claim during the prosecution of the '937 patent confirms this. For example, in the March 15, 2004 Response to Office Action, the applicant amended claim 1 to further clarify this point, adding “*the portion of the resources being specified as a quality of service guarantee*” and “in response to *the portion of the resources* allocated to the process being overloaded, increasing the quality of service guarantee for the process.” See Second Declaration of Jonathan R. DeBlois (hereinafter “2d DeBlois Decl.”)¹ Ex. R at 2. Giving key context to the amendment, the applicant stated that it was made because the examiner and applicant disagreed as to whether a resource could both meet a quality of service guarantee and be overloaded, as well as to clarify the term “quality of service guarantee.” *Id.* at 14.

The examiner, however, was not convinced, and in the April 21, 2005 Response to Office Action, the applicant canceled all claims that included “quality of service guarantee” being modified or increased as a result of an overload. Instead, the applicant relegated the phrase to the preamble of what would issue as claim 1, noting only that “resources allocated to the virtual server being specified as a quality of service guarantee.” In the body of the claim, in place of the language previously claiming an increase or modification of the “quality of service guarantee” to cure an overload, the applicant added the disputed term “modify a resource allocation for the virtual server.” *Id.* at 19-22. The examiner subsequently allowed the claims.

As can be seen by review of the entire prosecution history there was no disclaimer as to the resource allocation terms because the applicant’s arguments were directed to the scope of “quality of service guarantee.” Indeed, in every instance where the applicant made statements regarding “quality of service guarantee” the comments were focused on qualifying the phrase as specifying or representing a dynamic resource allocation. This makes logical sense, since as the

¹ All exhibit citations are to the Declaration of Jonathan R. DeBlois attached to IV’s Opening Claim Construction Brief unless otherwise noted.

examiner correctly noted, the term quality of service guarantee typically refers to a static allocation, not an allocation that can be dynamically modified. *See* VMware Br. at 4.

Furthermore, the fact that the applicant used the term “quality of service guarantee” in only three of the twenty-five claims of the ’726 patent supports IV’s proposed construction, not VMware’s argument. For instance—even setting aside the fact that the phrase is not used in any claims of the ’686 patent for the moment—its use in the preamble of claim 1 of the ’726 patent is simply to note that in the claimed system “the computer resources allocated to the virtual server [are] specified as a quality of service guarantee.” *See* Ex. F at claim 1. This is not saying that the resource allocations are quality of service guarantees, rather, one way in which a resource allocation can be represented is through a quality of service guarantee. This point is further highlighted by the applicant choosing to include the phrase in certain claims of the ’726 patent, while leaving it out of others, (and the ’686 patent), entirely. *See Inline Plastics Corp. v. EasyPak, LLC*, 799 F.3d 1364, 1371 (Fed. Cir. 2015) (error to construe disputed term to include an element expressly claimed in subsequent claim but not the claim at issue).

Finally, VMware’s argument that IV’s proposed construction reads out the term “allocation” is unfounded. As one of ordinary skill in the art would readily understand, for resources to be used by a virtual server those resources necessarily must be allocated. Therefore, IV’s proposed construction fully accounts for the disputed ‘allocation’ term. That being said, if the court is so inclined, IV would be amenable to removing the iterations of “each” in its proposed construction and replacing them with “the” to allay any potential concerns regarding its construction being unclear.²

B. “resource unavailable messages” / “denied requests to modify a resource allocation” (’686 patent claims 5-7)

IV’s Proposed Construction	VMware’s Proposed Construction
“an indication that a request by the virtual server cannot be immediately serviced” / “a request by	“indications that requests by the virtual server for additional resources are either implicitly or

² This would make IV’s construction “modif[y/ied] set of functions and features of a physical host used in implementing tasks for the virtual server.”

IV's Proposed Construction	VMware's Proposed Construction
the virtual server that cannot be immediately serviced”	explicitly denied, resulting from requests to modify a resource allocation” ³ see also construction of “modify a resource allocation”

VMware hinges much of its argument with respect to this term on its incorrect statement that “[b]oth parties have proposed this term for construction.” VMware Br. at 7. In actuality, VMware proposed the term “resource unavailable messages resulting from denied requests to modify a resource allocation” as a single term for construction, while IV proposed the two constituent terms “resource unavailable messages” and “denied requests to modify a resource allocation.” Clarifying which terms were actually proposed by which party renders VMware’s argument that IV’s construction attempts to read out the phrase “resulting from denied requests to modify a resource allocation” incorrect.

IV did not propose the entire phrase “resource unavailable messages resulting from denied requests to modify a resource allocation” for construction because the term ‘resulting from’ requires no construction. Thus, the dispute lies with the two terms connected by the ‘resulting from’ language, i.e. “resource unavailable messages” and “denied requests to modify a resource allocation” (each of which IV proposed to construe). With those two terms construed, it is clear that when they are connected using the phrase ‘resulting from’ that one is the result of the other. VMware itself acknowledges this in its proposed construction, which simply construes “resource unavailability messages” and then tacks on the remainder of the disputed term verbatim. *See* VMware Br. at 7 (“indications that requests by the virtual server for additional resources are either implicitly or explicitly denied, *resulting from denied requests to modify a resource allocation*”) (emphasis added).

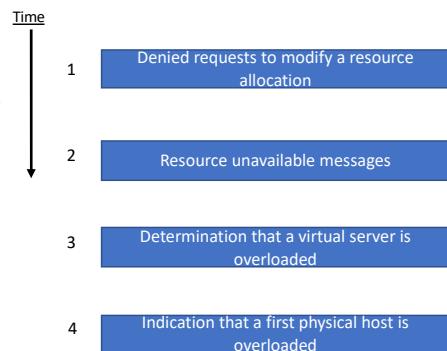
Neither is IV’s construction “circular” as VMware alleges. This becomes particularly evident when the disputed terms are read in the proper order, i.e., the order in which they occur within the computer system. For illustrative purposes the graphic below uses the relevant

³ VMware proposed this construction as a ‘compromise’ during the parties’ efforts to reduce the number of terms but did not indicate to IV that it would adopt that proposal as its new construction. As a result, IV’s Opening Brief included an older version of VMware’s proposed construction.

elements of claim 5 of the '686 patent as an example to show how the order of the elements is key to construing the disputed terms:

5. A method performed by a computing device, having a processor and memory, for modifying the computer resources allocated to a virtual server operating in a first physical host of multiple physical hosts, comprising:

receiving ⁴an indication that a first physical host is overloaded, wherein the indication is based on a determination that a virtual server is overloaded and wherein the ³determination that a virtual server is overloaded is based on one or more ²resource unavailable messages resulting from denied requests to modify a resource allocation;¹



As can be seen above, first there is *a denied request to modify a resource allocation*.

Subsequently, as a result of that denied request, there is *a resource unavailable message* generated. Based on the resource unavailable message, *a determination that a virtual server is overloaded* is then made. Finally, after it is determined that the virtual server is overloaded, *an indication that a first physical host is overloaded* is received. Importantly for the disputed terms here, the ‘resource unavailable messages’ are simply the result of the ‘denied request to modify a resource allocation,’ which are then used to make ‘the determination that a virtual server is overloaded.’

With this proper sequence in mind, the specification of the '686 patent describes how the two disputed terms should be construed. Because the claim language requires the ‘*determination that a virtual server is overloaded*’ be based on ‘*resource unavailable messages*,’ which themselves are representative of ‘*denied requests to modify a resource allocation*,’ the construction analysis begins with how the specification describes the ‘*determination that a virtual server is overloaded*.’ At column 2 lines 62-65 such a determination is disclosed as being based on the monitoring of ‘resource denials.’ The specification further defines the term ‘resource denials’ as “any request by the virtual server that cannot be immediately serviced.” Ex. E at 65-66. Therefore, it logically follows that the ‘*denied requests to modify a resource*

allocation' equate to 'resource denials,' and should be properly construed as "a request by the virtual server that cannot be immediately serviced." Following this same logic, the '*resource unavailable messages*' should be construed as "an indication that a request by the virtual server cannot be immediately serviced," since the indication is how the denials are quantitatively accounted for by the computer system.

This makes perfect sense when considering the operation of the claimed invention in its entirety and as reflected in the specification. The invention is a computer system that can only take action based on computer instruction sets. A denied request to modify a resource allocation, alone is not enough to cause the computer system to take action. Instead, in order for the system to react to the denied request, an instruction must cause the system to do so. Since in the claim, the denied requests must be quantitatively accounted for to make a determination of a virtual server overload, the computer is programmed to send a message that a request to modify a resource allocation has been denied (i.e., the claimed '*resource unavailable messages*'). *See Ex. E at 7:47-55.* These indications then become representative of the denied requests and are used as a quantitative measurement to determine whether the virtual server is overloaded. *Id.* Accordingly, IV's proposed construction is consistent with both the intrinsic record and the actual purpose and operation of the claimed inventions. *See Ex. E at column 8.*

In addition, despite VMware's conclusory argument to the contrary, IV's proposed constructions of the disputed terms are entirely consistent with its proposed construction of '*modified resource allocation*.' Common sense dictates that a 'denied request to modify a resource allocation' is quite different than a 'modified resource allocation.' A modified resource allocation is a resource allocation that has been modified, i.e., something tangible; a new resource allocation. A denied request to modify a resource allocation, on the other hand, is a denied request, i.e. a resource denial. There is no inconsistency between IV's proposed construction of the disputed terms and its construction of '*modified resource allocation*.' The terms are two distinct concepts that by their very nature cannot be distilled into a single construction.

C. “determination that a virtual server is overloaded” (’686 patent claims 5-7)

IV’s Proposed Construction	VMware’s Proposed Construction
Plain and ordinary meaning	“determination that an average number of resource denials for a virtual server is beyond a pre-configured threshold” See also construction of “resource denials”

VMware’s arguments with respect to this term fail to take into account the knowledge of one of skill in the art and the full intrinsic record. First, VMware claims that because the parties agreed on the construction of “indication that a first physical host is overloaded” the disputed term must be construed as something other than plain and ordinary meaning, since it is similarly a “technical term.” *See* VMware Br. at 9. The fact that the disputed term is “technical,” however, is ultimately irrelevant as to whether it should be construed as having its plain and ordinary meaning. The litmus test for claim construction is not whether a term is technical, but rather, whether it requires a specific construction because the patentee either made a clear disclaimer or acted as his own lexicographer. *See Thorner v. Sony Computer Entm’t Am. LLC*, 669 F.3d 1362, 1365 (Fed. Cir. 2012). Otherwise, the claim should be given its plain and ordinary meaning as understood by one of skill in the art at the time of the invention. *See id.* In this case, one of skill in the art would readily understand that the disputed term is being used in its well-understood and customary manner. *See* IV Br. at 18-20; *see also Liebel-Flarsheim Co. v. Medrad, Inc.*, 358 F.3d 898, 913 (Fed. Cir. 2004).

VMware’s next point that the disputed term and the ‘physical host overloaded’ term are different, and thus require different constructions is (i) irrelevant and (ii) in any event, consistent with IV’s position. Namely, that the term an ‘indication that a first physical host is overloaded,’ in light of the specification and claim elements, requires a construction other than plain and ordinary meaning, while “determination that a virtual server is overloaded does not. Turning to the construction at hand, a virtual server is a well-known and commonly used term among those of skill in the art and means exactly what it says; a virtualized server. *See* Declaration of Dr. Robert Akl, D.Sc. (hereinafter “Akl Decl.”) at ¶ 31. An overload in the context of a virtual server is equally well understood and is further explained by the language of the claim itself,

namely, that it is based on resource unavailability messages generated from denied requests to modify a resource allocation. *See id.* As described in detail in the previous section, the patent's specification further discloses exactly what is meant by resource denial and unavailability messages. *See supra* Section 1(B). Therefore, the disputed term has a well-understood plain and ordinary meaning and it should be construed as the same.

VMware's argument that the specification provides a "clear definition" of this term at column 5 line 29 is incorrect. Indeed, far from reciting a clear definition, column 5 line 29 describes "a flowchart of an embodiment" of this portion of the invention. Ex. E at 5:29. Reading a limitation into the claims from a preferred embodiment violates the most fundamental claim construction tenets. *See Comark Commc 'ns, Inc. v. Harris Corp.*, 156 F.3d 1182, 1186-87 (Fed. Cir. 1998). Moreover, this portion of the specification does not describe how a virtual server is determined to be overloaded. Rather, it is describing how it is determined whether or not a virtual server resource is overloaded. Specifically, VMware's purported support for its construction is discussing "whether a particular virtual server resource is overloaded," and states that to make a determination of the same it uses "a number of well-known techniques," finally noting that the determination is with respect to whether the virtual server is "overloaded for the corresponding resource." Ex. E at 5:42-43. The disclosure relied on by VMware, at best, only lends ambiguity to its proposed construction.

VMware alleges that its proposed construction is the only possible option that would not render the disputed term indefinite for lack of written description. VMware, however, fails to point to any evidence supporting this allegation. To the contrary, the intrinsic record is replete with evidence that the disputed term is being used in its plain and ordinary meaning. *See, e.g.*, Ex. E at Abstract, 2:62-65, 5:15-17, 8:5-55. Therefore, the disputed term is not indefinite and should be construed to have its plain and ordinary meaning.⁴

⁴ VMware's allegation that IV "refused to explain what it believes the plain and ordinary meaning" of the disputed term is or "otherwise explain why it disagrees with VMware's proposal," is simply wrong. IV explained to VMware that its proposed construction was reading limitations into the claim and thus improper. IV also explained, via

D. “virtual server” ('686 patent claims 5-7)

IV's Proposed Construction	VMware's Proposed Construction
Plain and ordinary meaning, in the alternative “a virtual machine that resides on a physical server and uses the physical server’s resources but has the appearance of being a separate dedicated machine”	“a process executing on a host computer that accepts communication requests, and that is capable of receiving a quality of service guarantee from a physical host”

VMware attempts to limit the common term in the art “virtual server” in two ways and neither has support in the intrinsic record. First, VMware tries to limit the patent’s virtual server to “a process executing on a host computer that accepts communication requests” by principally relying on a contorted reading of the unrelated ’051 patent. Second, VMware asks the Court to read a limitation about what the virtual server can do from a description of the preferred embodiments into the claim. That violates two claim construction canons.

IV’s proposal of plain and ordinary meaning should be adopted because the disputed term is a basic building block of virtual computing and has a well-known and understood meaning among those with skill in the art. *See Akl Decl. at ¶ 31.* Furthermore, nothing in the intrinsic record of any of the asserted patents rises to the level of a clear disclaimer of scope or indicates that the patentee acted as his own lexicographer. For example, the ’686 specification goes into great detail regarding the background of the invention and the field of technology, all of which supports IV’s proposed construction of plain and ordinary meaning. In doing so the patentee discusses the concept of virtual servers replacing physical servers which were owned by an ISP and leased in their entirety to a single customer. The patentee notes that many customers, however, do not want (or need) the processing power of an entire physical server. A better solution would be to use a single physical server to host multiple virtual servers that customers could customize to their needs. This clearly indicates that the patentee is using the term in its plain and ordinary meaning, i.e., a server that is virtualized. *See generally, Ex. E at column 1; see also iFly Holdings LLC v. Indoor Skydiving Germany GmbH, Case No. 2:14-cv-1080-JRG-*

phone and email conversations, that the term should be given its plain and ordinary meaning as understood by one of skill in the art because the intrinsic record shows this to be the case, and that by its very nature a plain and ordinary construction requires no further definition.

RSP, 2015 WL 9258264, at *19 (E.D. Tex. Dec. 18, 2015) (“in other words, the plain and ordinary meaning of the phrase ‘said chamber having a width’ is exactly what it plainly states, the chamber has a width.”).

On the other hand, the intrinsic record provides no support to limit “virtual server” with “a process executing on a host computer that accepts communication requests” The claims themselves and the intrinsic record surrounding the term contradict VMware’s proposed limitation. First of all, VMware fails to note that the prosecution history of the ’937 patent demonstrates that the disputed term was not meant to be limited to or defined as a “process.” Specifically, up until the October 5, 2004 Response to Office Action the claims included the term “process” and not “virtual server.” *See* 2d DeBlois Decl., Ex. R at 23-40. In the aforementioned office action response, however, **the applicant amended the claims to remove “process” and instead use “virtual server.”** *Id.* Had the applicant intended the disputed phrase to be limited to a process the claims would not have been amended to remove that exact term in favor of “virtual server.”

This is further supported by the specification of the ’686 patent. For example, the following excerpt is a quote from column 2 and is helpful in exposing VMware’s flawed construction:

Different customers have **different virtual server needs**. For example, a company **A providing large quantities of data and information to its employees and customers will want to ensure that its virtual servers are always available** to perform a large number of tasks. Company A may be willing to pay a premium for a guaranteed high quality of service, **with high server availability and large amounts of processing power** always on-call.

Ex. E at 2:1-8. The patentee is describing a virtual server just as anyone with skill in the art would expect and understand, not in terms of something that differs from or alters its plain and ordinary meaning. A customer needs a server to provide data and information to employees and customers; the server must always be available and have sufficient processing power to perform a large number of tasks. The server in this case simply happens to be virtual. What it is not, on the other hand, ‘is a process executing on a host computer’ as VMware suggests. A process

cannot have “high server availability” or “large amounts of processing power.” A process also cannot “provide large quantities of data and information.” Perhaps most importantly, there is nothing about a process that is virtual, in other words, every computer can (and does) run processes to implement all sorts of functions (including implementing a virtual server).

VMware’s construction interjects confusion and uncertainty into the claims and ignores that virtual server was specifically used in place of ‘process.’⁵

VMware’s second limitation, requiring a virtual server to be “capable of receiving a quality of service guarantee from a physical host,” lacks intrinsic support and is inconsistent within the context of the claims. VMware’s proposed limitation comes from a portion of the specification clearly labeled as describing embodiments. Ex. E at 49-51. VMware cites to one embodiment where virtual servers are capable of receiving quality of service guarantees and claims that the construction therefore must include that phrase. But VMware fails to note that immediately preceding that description the patentee states that “[r]eference will now be made in detail to several embodiments of the present inventions,” language clearly indicating that the invention is not limited to those embodiments. *Id.* Additionally, as discussed above *supra* Section 1(A) there are claims of the asserted patents that include the term ‘quality of service guarantee’ and claims that do not. Reading a limitation from a specific embodiment violates the most basic of claim construction principles. *See Comark Commc’ns, Inc.*, 156 F.3d at 1186-87.

E. “determining that a second physical host can accommodate the requested modified resource allocation” (’686 patent claims 5-7)

IV’s Proposed Construction	VMware’s Proposed Construction
Plain and ordinary meaning; in the alternative: “determining that a second physical host can accommodate the requests by the virtual server that could not be immediately serviced”	Indefinite, or in the alternative “determining that a second physical host can accommodate the denied request to modify a resource allocation”

The disputed term is not indefinite because it clearly refers back to the prior iteration of “denied requests to modify a resource allocation,” and therefore, has sufficient antecedent basis.

⁵ While IV’s alternate proposal more accurately captures the scope of the claims and spirit of the invention and should be adopted rather than VMware’s erroneous construction, however, IV’s position remains that a construction of plain and ordinary meaning is most appropriate here.

The requirement of antecedent basis is a rule of patent drafting. The MPEP states that a failure to include explicit antecedent basis is not always fatal to a claim. MPEP § 2173.05(e) (8th ed. Rev. 2 May 2004). The Federal Circuit has clarified that lack of antecedent basis does not render a claim indefinite “if the scope of a claim would be readily ascertainable by those of skill in the art.” *See Graphon Corp. v. Autotrader.com, Inc.*, Case No. 2:05-cv-530 (TJW), 2007 WL 1870622, at *11 (E.D. Tex. June 28, 2007) (citation omitted). When the instant claim is read in its entirety and in light of the specification the scope of the disputed term is readily apparent to one of skill in the art. For instance, the claim element immediately preceding the disputed term recites⁶, in relevant part, “resource unavailable messages resulting from denied requests to modify a resource allocation.” Furthermore, the specification clearly discloses what a denied request to modify a resource allocation is, how to monitor for it, and gives at least one embodiment complete with exemplary algorithms. *See, e.g.*, Ex. E at 11:12-57. One of skill in the art would clearly understand that the disputed term is referring back to the denied request to modify a resource allocation and therefore would also understand that the claimed second physical host would be required to have sufficient resources to satisfy the previously requested and denied resource allocation. *See* Akl Decl. at ¶ 32.

In addition, the same evidence supports IV’s position that the disputed term should be given its plain and ordinary meaning as opposed to VMware’s alternate proposed construction, which convolutes the term and, therefore, would be unhelpful to the jury. As discussed more fully in IV’s Opening Brief at Section B(7), there is a plain meaning in the art with respect to the disputed term’s constituent parts, and therefore, one of ordinary skill in the art would understand exactly what is meant by the disputed term. VMware’s proposal on the other hand, just inserts “the denied request to modify a resource allocation” in place of the disputed terms “the requested modified resource allocation.” It further adds “*see also* construction of “modify a resource allocation.” Not only is this unhelpful since it inserts one disputed term in place of another under

⁶ Claim 5 of the ‘686 patent is being used here as representative.

the guise trying to “remedy the antecedent basis problem,” it is unclear what VMware’s alternate proposed construction even is in light of the ambiguous “see also” citation it includes. IV’s alternate construction instead incorporates its construction of the term “denied request to modify a resource allocation,” so that its alternate proposal correctly reflects the intrinsic record and provides guidance to the factfinder as to the meaning of the disputed term. *See supra* Section 1(B).

2. Disputed Terms in the ’726 Patent⁷

A. “resource denials” (’726 patent claims 1, 4-5 & 8)

IV’s Proposed Construction	VMware’s Proposed Construction
“indications that requests by the virtual server cannot be immediately serviced”	“indications that requests by the virtual server for additional resources are either implicitly or explicitly denied”

As an initial matter VMware notes that the parties have agreed on the first part of the proposed construction—“indications that requests by the virtual server”—for the disputed term despite the fact that IV’s original proposal did not pluralize “indication.” Prior to VMware’s Opening Brief, IV was unaware that VMware was amenable to agreeing on this portion of the construction since VMware made no mention of it during the meet and confer process or the parties’ follow-ups. Nevertheless, IV agrees to modify its proposed construction to reflect this agreement.

The remainder of VMware’s proposed construction, however, should be rejected because, as discussed in IV’s Opening Brief, Section B(3), it improperly limits the disputed term to one narrow embodiment despite the intrinsic record defining the term more broadly. VMware draws its proposed construction from a cite to the specification that describes “instances wherein a request for additional resources is either implicitly or explicitly denied.” *See* Ex. F at 7:51-54. This cite is in the context of discussing a specific embodiment where the virtual server resource monitor monitors different types of resource denials. *Id.* at 5:41-43. The specification then gives

⁷ The ’726 patent terms that overlap with ’686 terms discussed above can be found in Exhibit 7 to VMware’s Opening Brief. These claim terms should be construed as IV proposes for the reasons stated above with respect to the corresponding ’686 patent term.

a description of that embodiment’s ‘resource denials’ which VMware adopts as it proposed construction. While an “explicit” or “implicit” denial can be types of resource denials, it is improper to limit the disputed term to this narrow exemplary embodiment, particularly in light of the broader disclosure that captures the term’s full scope. *See Comark Commc’ns, Inc.*, 156 F.3d at 1186-87. Indeed, IV’s proposed construction is drawn from the patent’s Summary of Invention that provides a broader disclosure of “resource denials.” Ex. F at 2:55-61. As such, VMware is not only asking the Court to read in a limitation, it is also seeking to read out explicitly disclosed claim scope. *See Duncan Parking Techs, Inc. v. IPS Grp., Inc.*, 914 F.3d 1347, 1362 (Fed. Cir. 2019).

B. “quality of service guarantee” (*'726 patent claims 1 & 4*)

IV’s Proposed Construction	VMware’s Proposed Construction
“a guaranteed resource allotment which can be dynamically increased/modified”	“information that specifies a guaranteed amount of an assigned resource, and that can be dynamically increased/modified”

VMware argues that IV’s construction provides no clarity “on what is being guaranteed” despite the plain language of IV’s proposal specifying that an “amount of an assigned resource” is the item that is guaranteed. IV’s construction is consistent with the specification’s disclosure of how resource allocations are linked to “quality of service guarantee.” IV’s proposed construction diverges from the specification only by using the synonym “allotment” for “allocation.” Since “allotment” is synonymous with “allocation”⁸ it follows that IV’s proposed construction of “a guaranteed resource allotment which can be dynamically increased/modified” is much truer to the specification than VMware’s proposal. The specification illustrates this point in the background section when noting “[h]owever, customers may be unable to anticipate the exact amount of resources they will require, and a static **assignment of a particular resource allocation** limit may not allow the virtual server system to adapt to changing customer needs . . . a better **resource allocation model** . . .” *See* Ex. F at 2:14-19. This passage both links the term

⁸ IV would be willing to change its proposed construction to use “allocation” in place of “allotment” to avoid any potential confusion, if the Court is so inclined.

‘assignment’ to something that can’t be altered, and also indicates an ‘assignment’ is distinct from the phrase ‘resource allocation.’

On the other hand, VMware’s proposal of ‘information that specifies a guaranteed amount of an assigned resource’ mixes terms and concepts from other unrelated parts of the patent into “quality of service guarantee.” For instance, VMware’s construction inserts “information that specifies” and “an assigned resource” both of which interject ambiguity into the disputed term in light of the specification’s disclosure. ‘Quality of service guarantee information’ is described in the specification as distinct from a ‘quality of service guarantee.’ *See Ex. F at 4:39-60.* The former being how a quality of service guarantee is represented in some embodiments, such that the information can be stored in a “parameter table,” while the latter is the actual resource allotment that the virtual server is guaranteed. *Id.* at 2:38-46. Furthermore, the phrase ‘assigned’ appears only twice in the specification both in the context of criticizing the prior art as teaching “fixed” or “static” assignment of resources. *See id.* at 2:5-6, 2:14-16. This contradicts the patented invention’s teaching of “dynamic” resource allocation and VMware’s own proposed construction.

3. Alleged Means Plus Function Terms for the ’726 and ’686 Patents

A. ’686 patent claim 7 “component” terms (i.e., clauses 1-3 of Ex. A)

As discussed in detail in IV’s Opening Brief, Section B(10)(i), the ‘component’ claims of the ’686 patent (listed in full in Ex. A, clauses 1-3) are not subject to § 112 ¶ 6 because: (1) “means for” is not present in any claims, (2) the prefatory language used does not consist of merely a nonce word followed by function, and (3) the claim language itself provides sufficient structure to avoid the application of § 112 ¶ 6. *See Fisher-Rosemount Sys. v. ABB Ltd.*, Case No. 4:18-cv-00178, 2019 WL 6830806, at *15-*16 (S.D. Tex. Dec. 12, 2019). IV will not re-hash these arguments herein in the interest of brevity. In its Opening Brief, however, with respect to Clauses 1-2, VMware argues that the ’686 specification fails to disclose any algorithm that performs the claimed functions. Assuming that the terms are found to be subject to § 112 ¶ 6—

which they should not be—the specification does in fact disclose corresponding structure and algorithm for these disputed terms.

Clause 1, which claims the function of “receiv[ing] an indication that a first physical host is overloaded,”⁹ has sufficiently disclosed structure in the intrinsic record, including within the claim itself. This fact would be easily understood by one of skill in the art, which is the standard for the sufficiency of structure in a § 112 ¶ 6 analysis. *See Telcordia Tech., Inc. v. Cisco Sys., Inc.*, 612 F.3d 1365, 1377 (Fed. Cir. 2010) (“the specification need only disclose adequate defining structure to render the bounds of the claim understandable to an ordinary artisan.”).

First, the claim elements preceding the “component” element at issue here provide structure for this function, specifying that the ‘component configured to receive’ is within a physical host that includes one or more processors and one or more memories. These are structural terms that in combination with further description in the specification—as discussed below—clearly disclose the corresponding structure. The ’686 patent additionally discloses structure corresponding to the claimed function in Figure 1, 2A, 2B, and 5, and the related descriptions of the same. Specifically, Figure 1 shows that the invention includes a “dynamic resource configuration module 100” which comprises, among other things, a “physical host load balancer 130.” Figure 5 shows that within the physical host load balancer 130 there is a “physical host resource monitor 540 and load balancing calculator 530.” The specification discloses that:

physical host load balancing module 130 receives a signal 510 from the virtual server resource modifier 120 indicating that virtual server 162B requires an increased resource allocation. This signal is used as an input 520 into the load-balancing calculator 530. The load-balancing calculator 530 also requests and receives as input the current physical host resource loads 535 from the physical host resource monitor 540. The physical host resource monitor 540 performs periodic physical host resource checks 545 upon the group of physical hosts . . . the load-balancing calculator 530 determines whether physical host 160A is capable of supporting the request for additional virtual server 162B resource 510. If the

⁹ VMware claims that the entire element should be considered the function, however, the portion of the limitation following IV’s proposed function is merely further description of the claimed “indication,” not the function itself which is “receiving an indication that a first physical host is overloaded.” The presence—or lack thereof—of the specific properties of the indication is an infringement/non-infringement issue, not a claim construction issue.

resource request exceeds the available resources of physical host 160A, the load-balancing calculator 530 determines that physical host 160A is overloaded.

Ex. E at 11:4-26. (emphasis added). Accordingly, the load-balancing calculator 530, which is within the physical host load balancing module 130 contained within the dynamic resource configuration module 100, is the structure that receives the indication that a first physical host is overloaded. The indication itself is the current physical host load on 160A received from the physical host resource monitor 540 compared to the request for additional resources 510 by virtual server 162B. Combined with the flow diagrams in Figure 2A and 2B, as well as the structural limitations of the claim itself, this provides sufficient disclosure to one of skill in the art of what structure performs the claimed function. *See Akl Decl. at ¶ 33.*

Clause 2 finds similar structural support in the claim and throughout the specification for the function of “determin[ing] that a second physical host can accommodate the requested modified resource allocation.” For example, the starting point is again within claim 7 itself which specifies that the function is performed within a physical host server with one or more memories and one or more processors, as discussed above. Figure 1, Figure 5, and the specification portions in which they are discussed provide additional structure, as also discussed above. Furthermore, picking up where the above block quote ended, the specification describes in detail how the physical host load balancer 130 and the load-balancing calculator 530 perform the claimed function:

[T]he load-balancing calculator 530 uses an easiest fit heuristic to find the physical host that has the most available resources. Each different type of resource is associated with an ordinal and a weight . . . [discloses easiest fit heuristic algorithm]¹⁰ . . . [u]sing the easiest fit heuristic, the load-balancing calculator 530 will select the physical host with the largest weighted resource availability to receive the virtual server 162B (in the above example of Fig. 5, physical host 160B) . . . the load balancing calculator 530 sends 550 a signal 560 to the dynamic virtual server mover 140 indicating that the virtual server 162B is to be transferred to physical host 160B.

Ex. E at 11:27-57 (emphasis added). Therefore, the structure corresponding to the claimed function is set out in detail and includes, the load-balancing calculator 530, which is within the

¹⁰ For the specific ‘easiest fit heuristic’ formula *see* ‘686 Specification at 11:30-45.

physical host load balancing module 130 contained within the dynamic resource configuration module 100. When viewed in light of the easiest fit heuristic formula, the flow diagrams in Figure 2A and 2B, as well as the structural limitations of the claim itself, this provides sufficient disclosure to one of skill in the art of what structure performs the claimed function. *See* Dr. Akl Decl. at ¶ 34.

With respect to Clause 3, IV agrees with VMware's identified function and a portion of the corresponding structure that it has identified, however, VMware's structure is too narrow and leaves out key structural elements that play a part in performing the claimed function. For example, VMware argues that virtual server mover 140 as shown in Figure 6 is the entirety of the corresponding structure for the function. VMware's structure proposal, however, fails to account for the 'generate a physical host transfer signal' portion of the claimed function. The virtual server mover 140 only performs the transfer of the virtual machine from a first to a second host, it does not perform the generating of the transfer signal. *See* Ex. E at 11:27-57.

Therefore, at least the load balancing module 130 and load balancing calculator 530, which are components of dynamic resource configuration module 100, perform the 'generate' function. *See id.* at 11:54-57; Fig. 5. Furthermore, in order to generate the transfer signal, the physical host load balancer 130 needs to receive information from the virtual server resource monitor 120 relating to virtual server 162B's resource request, as well as information from the physical host resource monitor 540 relating to the resource loads on all physical hosts 160A-C. All of these components are located within dynamic resource configuration module 100, and should all be included in the corresponding structure. *See* Akl Decl. at ¶ 35.

B. '726 patent claim 1, 4 & 5 (i.e., clauses 4-8 of Ex. A)

As an initial matter, the claim elements reflected in Clauses 4-8 of Ex. A should not be subject to § 112 ¶ 6 for the reasons set forth in IV's Opening Brief, Section 10(ii)-(vi), namely, because one of ordinary skill in the art would be well-informed of the structure of the elements from the claim language and specification disclosures. *See* Akl Decl. at ¶ 36. Should the Court

disagree and find that § 112 ¶ 6 does apply, however, each disputed element has sufficient corresponding structure identified either in the claim itself, specification, or both. *Id.*

Additionally, VMware attempts to read the term “creating” into each one of its proposed functions, despite the fact that only claim 4 actually claims “program code for creating,” while claims 1 and 5 make no mention of the term. VMware should not be allowed to limit the function of an alleged means-plus-function element to a function—“creating”—that is not found in the disputed term. Therefore, IV will address the iterations of the disputed terms of claims 1 and 5 herein. As for the disputed “program code” elements of claim 4, IV agrees with VMware’s proposed function to the extent it contains “creating,” but disagrees with VMware’s proposed structure for the same reasons discussed above with respect to the non-program code elements.

Finally, as with the alleged means-plus-function elements in the ‘686 patent discussed above, VMware seems to misapprehend the scope of what is considered an “algorithm.” For purposes of § 112 ¶ 6, when an algorithm is required to support a function (which it is not here), that algorithm does not have to take the form of a mathematical formula. Anything adequate to permit one of ordinary skill in the art to perceive the bounds of the invention is sufficient, and can be presented in the specification read as a whole in any understandable terms including as a mathematical formula, in prose, or as a flow chart, or in any other manner that provides sufficient structure. *See Intelligent Automation Design, LLC*, 2020 WL 486830, at *3-*5.

i. “virtual server resource monitor” term, clause 4 of Ex. A

The term ‘virtual server resource monitor’ as claimed in claims 1 and 5 is required to, (a) monitor resource denials, and (b) send a virtual server overload signal in response. The claim itself provides that this takes place in a physical host server including a processor and memory. Contrary to VMware’s conclusory assertions, ‘server’ can provide structure and has a well-known meaning in the art. *See Sound View Innovations, LLC v. Facebook, Inc.*, Case No. 16-cv-116 (RGA), 2017 WL 2221177, at *10-*11 (D. Del. May 19, 2017). Furthermore, the specification and figures provide additional structure beyond simply just virtual server resource

monitor 110, as proposed by VMware. The summary of the invention notes that “a dynamic resource configuration module monitors resource denials received by the virtual servers and determines if a virtual server is overloaded based upon the resource denials.” Ex. F at 2:52-55. Figure 1 depicts the dynamic resource configuration module 100 as including virtual server resource monitor 110, virtual server resource modifier 120, physical host load balancer 130 and dynamic virtual server mover 140. Figure 3, which also depicts the virtual server resource monitor 110, further includes virtual server 162B, recognizing that the virtual servers and physical hosts also necessarily make up part of the structure performing the claimed function.

Additionally, the specific description of VMware’s proposed “algorithm” is too narrow and adopting it would impermissibly narrow the term. For example, one of ordinary skill in the art would know from the claims and specification disclosure that there are many different ways to intercept systems calls, and resource denials can be accounted for and an overload determined in ways other than storing the denials in individual tables and using a pre-specified time window. *See Akl Decl.* at ¶ 37. The patentee recognizes this in several places throughout the specification. *See, e.g.,* Ex. F at 8:46-49; 8:65-66; 9:43-46. So, while IV does not dispute that virtual server resource modifier 110 is a part of the corresponding structure, the Court should also include the other components noted above, particularly dynamic resource modification module 100, and disregard VMware’s narrow recitation of the “algorithm.”

ii. “virtual server resource modifier” term, clause 5 of Ex. A

VMware again points to an overly narrow description of the structure performing the claimed function. The same arguments from Section (i) above apply to this disputed term. The virtual server resource modifier is certainly part of the disclosed structure, but the Court should also include dynamic resource configuration module 100 and the virtual servers in order to capture the full scope of the corresponding structure. One of ordinary skill in the art would understand that dynamic resource configuration module 100 is not a “black box” as VMware

contends, but rather, is a well-known component that plays a part in performing the claimed function.¹¹ *See Akl Decl.* at ¶ 38.

iii. “load balancing module” term, clause 6 of Ex. A

As with the other terms discussed above in sub-sections (i)-(ii), VMware’s proposal for this disputed term includes part of the structure that performs the claimed function but leaves out others. It ignores the fact that physical hosts 160A-C and virtual servers 162A-C are described as contributing to the claimed function (as seen in Figure 5 and the related description discussed above in Section 3(A)), as well as the fact that the physical host load balancing module is part of the dynamic resource configuration module 100, and includes the physical host resource monitor 540 and load balancing calculator 530, all of which should properly be considered as part of the corresponding structure. *See Akl Decl.* at ¶ 39.

iv. “dynamic virtual server mover” terms, clauses 7 & 8 of Ex. A

The corresponding structure for this disputed term is similarly broader than VMware’s description and many of the same arguments from Section 3(A) and sub-sections (i)-(iii) above are applicable here as well. Primarily among those arguments, VMware ignores the physical servers, virtual machines and dynamic resource configuration module 100, each of which are structural elements that play a part in performing the claimed function. *See supra* Section 3(A). For instance, the virtual server mover is part of the dynamic resource configuration module which is coupled to the physical hosts running the virtual machines. *See Ex. F at Fig. 1 & 5; see also, Ex. F at 2:52-67.* In addition, the specification describes physical host load balancer 130 as using the dynamic virtual server mover 140 to transfer virtual servers between hosts, indicating that the load balancer (also part of dynamic resource configuration module 100) is used, in part, to perform the claimed function. *Id.* at 5:9-13. Finally, one of skill in the art would understand from the specification’s disclosures that the ‘make, then break’ and ‘break, then make’ methods

¹¹ It’s also worth noting that as with all of its proposed constructions, VMware seems to incorporate its constructions of other terms, in this case, “virtual server overload signal,” “modify a resource allocation,” and “virtual server,” into the construction of the disputed term. This makes its proposed construction even more confusing, but most importantly, shows that it believes the structure of the disputed term at least includes the above identified terms, which belies its arguments in favor its narrowly defined structure.

of transfer are not the only methods for transferring virtual servers and thus the structure should not be so limited. *See Akl Decl.* at ¶ 40.

The '752 Patent

United States Patent No. 7,949,752 (“the ‘752 patent”) relates to the use and customization of network systems that allows delivery of highly reliable applications and services from the cloud to remote users. Through the use of novel virtualization systems and methods a client machine need only have basic browser support in order to have full use of cloud-based programs and applications. The network system and methods make this possible through virtualizing physical hardware into logical abstractions allowing resources to be distributed and managed at locations remote from end-user clients.

4. Disputed Terms in the '752 Patent

A. “exhausted” ('752 patent claims 1, 9 and 24)

IV’s Proposed Construction	VMware’s Proposed Construction
“used up to the allotted or pre-determined amount”	“unavailable for reuse”

VMware claims that during the prosecution of the '752 patent the patentee made a clear disclaimer regarding the scope of the disputed term, and therefore, the term should be construed according to its proposal of ‘unavailable for reuse.’’ VMware then argues that in construing the term as IV proposes in prior litigation Judge Mitchell committed an error and misapplied the *Ecolab* case. Neither assertion is correct.

As Judge Mitchell found in the HCC Case, the prosecution history, read as a whole, does not contain a clear and unambiguous disclaimer with respect to the disputed term. While in response to a Section 102 rejection the patentee did state that the prior art failed to disclose a resource that is exhausted upon being consumed because it “teaches the exact opposite—i.e. that the ‘cartridges’ in Chou are reused,” the examiner disagreed, calling the argument “misleading” and further noting that “in the applicants invention, resources are reused.” *See* 2d DeBlois Decl., Ex. S at 13. The examiner then pointed the applicant to the exact portion of his specification that discloses reuse. *See* 2d DeBlois Decl., Ex. S at 19 (“On page 25 of the applicant’s specification

resources are defined as including processing time, memory storage space, and the like. While these resources may be at full capacity at any given time, the[y] can always be reused later they are not at full capacity.”). VMware misinterprets this as a mere ‘disagreement’ between the patentee and the examiner, noting that in the responsive office action the examiner withdrew the Chou reference and put forth another piece of prior art, “Cohn,” that formed the basis of his “new” Section 102 rejection. VMware then claims that Cohn disclosed “consuming non-reusable fees for using a service,” which proves that there was a disclaimer and the examiner agreed, at least passively. This is not entirely accurate, however, because VMware fails to mention two key facts; (1) the Chou reference was not withdrawn because the examiner somehow acquiesced to the patentee’s statement, and (2) the Cohn reference discloses reusable service resources as well.

With respect to the first issue, a careful analysis of the prosecution history reveals the flaw in VMware’s argument. In the same OA response where the patentee made the alleged disclaimer the applicant also amended then claim 86 to include the term “an amount of” the service resource being exhausted upon being consumed. *See* 2d DeBlois Decl., Ex. S at 2. Because Chou disclosed software containers called “cartridges” that the examiner equated to the claimed service resource, an ‘amount of’ the cartridge could never be exhausted regardless of the scope of the term ‘exhausted.’ Thus, the examiner’s withdrawal of Chou as a Section 102 reference was unrelated to the meaning of the disputed term.¹²

Regarding the second issue, a careful reading of the Cohn reference and file history shows that the reference also disclosed service resources that could be reused. In Cohn the examiner considered the service resource the “information providers” that comprise databases of information users can access and then be charged for based on level of access and the like. This is clearly a re-usable resource. Consider a scenario where subscriber A gets access to the database for information Y, information Y is not then “unavailable for reuse,” subscriber B can

¹² This makes logical sense considering that had the examiner truly believed that the patentee’s invention claimed the re-use of service resources—which he clearly did—as a matter of law he could not have withdrawn the objection and selected a reference that disclosed only non-reusable service resources simply to avoid the issue.

also subscribe to and receive access to information Y. VMware’s claim that the Cohen reference disclosed “consuming non-reusable fees for using a service” thus evidencing a disclaimer in scope of the disputed term, is therefore incorrect.¹³

When the prosecution history and the references discussed therein are properly analyzed, it becomes apparent that, just as Judge Mitchell concluded, there was no clear disclaimer with respect to the disputed term. In fact, the evidence shows that the patentee recognized that the examiner was correct in his assertion that “exhausted” is broader than “unavailable for reuse” by abandoning that argument despite the fact that the Cohn reference disclosed reusable service resources. Instead, he amended the independent claims as suggested by the examiner to include the “URL” elements found in dependent claims. *Compare* 2d DeBlois Decl., Ex. S at 19-23 *with* Ex. S at 24-35. Therefore, VMware has not (and cannot) meet its burden of showing that there was a clear and unmistakable disclaimer to the scope of the disputed term.

VMware’s dismissal of Judge Mitchell’s logic in coming to the same conclusion regarding whether a disclaimer occurred is equally as flawed as its other arguments noted above. For example, VMware claims that Judge Mitchell committed an error in applying the *EcoLab* case. Judge Mitchell, however, considered this very same argument when it was made by HCC and specifically took note of it in her order, ultimately concluding that “[w]hile the facts in *EcoLab* may be distinguishable from the current case, its holding is nevertheless instructive. The prosecution history, when considered as a whole, demonstrates that the patentee did not make a clear and unmistakable disclaimer.” *See* Ex. N at 33. Accordingly, the intrinsic record, a Primary Examiner at the USPTO, and a United States Magistrate Judge rejected VMware’s position as should this Court.

¹³ This characterization is also illogical because even if the reference was disclosing something similar to “consuming long-distance minutes” which VMware analogizes to, the service is the ability to make a long distance phone call and the service resource is the phone lines, switches, etc. that carry the phone call. Thus, even in VMware’s characterization the service resource could never be “unavailable for reuse” since the phone lines and switches used to make long distance phone calls (regardless of the pricing structure for customers) are reused countless times.

B. “consumed” (’752 patent claims 1, 9 and 24)

IV’s Proposed Construction	VMware’s Proposed Construction
“used”	“used up”

VMware’s arguments as to why the disputed term should be limited to “used up” have no factual or legal basis. As an initial matter, VMware’s arguments here are a transparent attempt to save its carefully crafted non-infringement position regarding the ‘exhausted’ term previously discussed should the Court not find in its favor there. In addition, the idea that the disputed term was somehow “narrowed” during prosecution is equally unfounded. The claim amendment that VMware points to as support was made to specify that the agent is doing the consuming, not that “use” and “consumed” have two separate and distinct meanings. In fact, the portion of the file history that VMware cites to as support shows that the amendment was unrelated to the definition of the “use” or “consumed” terms, but rather, was made to specify that the agent is not a GUI as the examiner contended because a GUI cannot use the hardware resources of a server device. *See* VMware Br. Ex. 15 at 10; *see also* 2d DeBlois Decl., Ex. S at 38-40.

Furthermore, VMware’s argument that somehow the fact that ‘consumed’ is used in the past tense (as opposed to consume) indicates that it means “used up” overlooks the obvious point that the claim states “to be consumed,” which is the future tense. Thus, nothing has been used up, rather, something will be used. Similarly, VMware’s argument that there is a “definitional statement” in the specification defining ‘consumed’ as ‘used up’ also fails. Such a claim is belied by the fact that while the disputed term is referenced in a single instance as ‘used up’ it is referred to interchangeably with ‘used’ at least half a dozen times. *See* IV Br. at 5.

Finally, even the examples cited by VMware regarding voice mail and call placement are misplaced. VMware is correct if it is implying that the time a human spends making a phone call or leaving a voicemail cannot be reused, i.e., one cannot literally travel back in time, however, this is not what the patent is referring to. When the patent speaks of discrete units of a service resource being consumed while using a voicemail or long-distance telephone service the point is just that in one embodiment service resources can be structured in such a way that allows a

service provider to charge for each unit used. Either way, in no instance is the term limited to “used up,” a fact recognized by Judge Mitchell when facing the same arguments VMware makes now. *See Ex. N at 29-31.* Therefore, the Court should disregard VMware’s attempts to narrow the disputed term and adopt IV’s construction which is overwhelmingly supported by the intrinsic record.

C. “service” (*’752 patent claims 1, 3, 9 and 24*)

IV’s Proposed Construction	VMware’s Proposed Construction
“network functionality available to agents”	“an application that is used on behalf of a principal”

The Court should adopt IV’s proposed construction for this term because it is taken directly from the extensive specification disclosures, whereas VMware points to a single instance in which the disputed term is referenced as it proposes. There is no dispute that the patent discloses that a “service” could be compiled into a discrete program which the claimed agent could use based on an end-user’s (i.e. principal’s) instruction. *See* IV Br. at 7. But this is merely a single embodiment. Limiting the disputed term to one narrow embodiment mentioned a single time in the 35 pages of specification disclosure is inconsistent with basic claim construction principals. *See Comark Commc’ns, Inc.*, 156 F.3d at 1186-87. In this case, the overwhelming weight of the remaining disclosures indicate that the disputed term should be construed as IV proposes. For example, the specification notes that “an agent is operable to use a service.” *See* Ex. D at abstract. Examples of services include an email service, a voice mail service, etc., each supported by a “sub-system” of the network, i.e. network functionality, not applications. *Id.* at 13:21-28, Fig. 2. Further, the result of using a service might benefit a human user, but they are being used by the claimed agent. Therefore, the Court should adopt IV’s proposal because it captures the full scope of the disputed term as used in the intrinsic record and the claims.

D. Means-Plus-Function Terms

The disputed terms, as well as the Parties’ proposed function and structure can be found in Ex. B. The main dispute surrounds VMware’s attempts to limit its proposed structure by qualifying its identification with the phrase “as described in,” followed by citations to a narrow

embodiment of such structure. For the reasons set forth in IV's Opening Brief, Section 3(A)(4), this should not be allowed.

VMware argues that § 112 ¶ 6 requires that it include specific specification citations that correspond to the disclosed structure, but this is not so. This Court has previously construed means plus function terms where the corresponding structure for performing the claimed function includes only the disclosed structure, e.g., 'widget 26,' and § 112 ¶ 6 requires nothing more. *See, e.g., VLSI Tech. LLC, v. Intel Corp.*, Case No. 1:19-CV-00977-ADA, Claim Construction Order, Dkt. No. 101. In any case, the real issue here is that VMware's citations exclude disclosed embodiments specifically linked to performing the claimed function and are unnecessary. Accordingly, the Court should adopt the proposed structure that the Parties agree on and not read in VMware's narrowing citations.

The only other dispute regarding the '752 patent's means plus function terms is with respect to the structure of "means for monitoring an amount of the service resource used by the network-based agent." The Parties' proposals are below:

IV's Proposed Construction	VMware's Proposed Construction
<u>Function</u> : monitoring an amount of the service resource used by the network-based agent" <u>Structure</u> : Service Wrapper 26	<u>Function</u> : monitoring an amount of the service resource used by the network-based agent" <u>Structure</u> : monitor as described in '752 patent at 16:50-61.

The Court should adopt IV's proposed structure because it captures the full scope of the specifications disclosure whereas VMware's proposal includes only a sub-component of the service wrapper 26 which performs the function as a whole. *See* IV Br. at 8. Furthermore, VMware's argument that the converter sub-component of service wrapper 26 is irrelevant to performing the claimed function is misplaced. Since the converter converts the computer language used by a service to the computer language used by the agent server, which in turn interacts with monitor 50 to enforce service resource usage of the agents when using the service, the converter does perform part of the claimed function. In other words, without the converter it would not be possible to monitor and enforce service resource usage since the components using the services and service resources would be unable to communicate with the components

monitoring the same. *See* Ex. D at 16:22-61. Therefore, the Court should adopt IV’s construction as it accurately reflects the scope of the corresponding structure.

The ’051 Patent

The ’051 patent teaches combinations of virtualized components, tunnels, and physical interfaces that collectively allow enterprise or data center operations to be siloed on a per-tenant basis so that each tenant may use the same physical hardware as another tenant without risk of security breach or hardware failure. And, because the tenant’s servers are virtualized, the addresses of one tenant’s virtual servers may overlap with the addresses of another tenant’s virtual servers. Such virtual isolation thus reduces the risk of a security breach, while continuing to provide tenants with a lot IP addresses.

5. Disputed Terms in the ’051 Patent

A. “virtual server” (claims 1, 3, and 6)

IV’s Proposed Construction	VMware’s Proposed Construction
Plain and ordinary meaning Alternatively – “virtual machine(s) that reside(s) on a physical server and use(s) the physical server’s resources, but [has/have] the appearance of being a separate, dedicated machine(s)”	“a process executing on a host computer that accepts communications requests”

Notably, in the four pages dedicated to this term in its Opening Brief, VMware does not discuss IV’s primary proposal, that “virtual server” should be construed with its plain and ordinary meaning. As explained in IV’s Opening Brief, plain and ordinary meaning is the correct construction because the ’051 patent uses this well-understood term in a non-limiting manner. IV Br. at 41-42. At the time of the inventions of the ’051 patent, a POSITA would readily understand that the term meant a virtualized server and no further definition is necessary. *See* Declaration of Dr. Vijay Madisetti (hereinafter “Madisetti Decl.”) at ¶¶ 31-33. As such, “virtual server” as used in the patent, has an even broader meaning than even IV’s alternative construction. IV’s alternative proposed construction, the entire focus of VMware’s Opening Brief, is meant to define the term in the context of known virtualization practices. In no event should the term be defined as VMware proposes. VMware’s construction finds no support in the

intrinsic record and runs counter to how a POSITA would have understood the term at the time of the invention.

Finding no support in the '051 patent's written description, VMware resorts to citations from a patent application incorporated by reference into the '051 patent's specification, U.S. Patent App. Ser. No. 09/432,286 ("the '286 application"). Not only does the '286 application's specification not define, or otherwise disclaim the scope of, virtual server in the manner advocated by VMware, but the citations relied on by VMware are describing a physical host server, not a virtual server. *See* VMware Br. at 34-38 (citing 6,976,258) (distinguishing "a server, which is a process, executing on a dedicated physical host" from "a virtual host server [that] services numerous client requests for multiple virtual hosts" in same paragraph); Abstract (differentiating "plurality of virtual hosts . . ." from ". . . on a single physical host computer"); 1:16-18 (referencing "a server program executing on a single physical host computer" indicating that servers, virtual or otherwise, are not exclusively referred to using "process" language). In addition to there being no teaching or suggestion in the '286 application that the "process" characterization of a server was intended to be limiting, there is no disclosure in the '051 patent purporting to limit the meaning of virtual server to particular examples or embodiments discussed in the '286 application. *See* Ex. H at 3:64-67 (characterizing incorporated reference as example disclosing "a method for creating such a private virtual server", but not expressly defining a virtual server as such).

Further, VMware's proposed construction fails to account for the "virtual" aspect of the disputed "virtual server" term. As discussed in IV's opening brief, VMware's proposed construction of "a process executing on a host computer that accepts communications requests" does not contain any wording that would enable a POSITA to distinguish the claimed virtual server from a networking-related process associated with a conventional, physical server.¹⁴ It is well-established that claim construction "begins and ends in all cases with the actual words of

¹⁴ Indeed, this is precisely because the "process" language from '286 application that VMware has latched onto was language describing an example of a physical, non-virtual server.

the claim.” *Renishaw PLC v. Marposs Societa' per Azioni*, 158 F.3d 1243, 1248 (Fed. Cir. 1998). VMware’s proposed construction violates this fundamental principle by disregarding the term “virtual.”

Although IV believes a construction beyond plain and ordinary meaning is unnecessary, the proposed alternate construction seeks to clarify the “virtual” aspect of a “virtual server” by way of an embodiment of a virtual server pertinent to this case. The specification makes it clear that a virtual server “can service requests to multiple network addresses or domain names [t]hus, the functionality of numerous hosts is provided by a single physical host computer” and further describes the inability of existing virtual servers “to allocate appropriate amounts of computer resources of the physical host computer to servicing client requests made to specific virtual hosts.” Ex. H at 2:50-64, 3:57-64. In light of such disclosure in the specification and its known meaning in the art, a POSITA would readily have understood a virtual server as an example of virtual machine that abstracts physical server resources and appears from the outside to be a fully-functional, dedicated physical server. *Id.* at 2:60-62 (“. . . the virtual host server can service requests to multiple network addresses or domain names. Thus, the functionality of numerous hosts is provided by a single physical host computer.”). *See Medrad, Inc. v. MRI Devices Corp.*, 401 F.3d 1313, 1319 (Fed. Cir. 2005) (it is “entirely proper to consider the functions of an invention in seeking to determine the meaning of particular claim language”); *see also* Madisetti Decl. at ¶¶ 31-33.

B. “physical interface[s]” (claims 1 and 3)

IV’s Proposed Construction	VMware’s Proposed Construction
Plain and ordinary meaning	“hardware that provides a point of communication between two or more devices”

The disputed term “physical interface” should be construed as its plain and ordinary meaning because it is a well-understood construct to those of skill in the art and requires no construction other than plain and ordinary meaning. The only exceptions to giving a term its plain and ordinary meaning are when the patentee has acted as his own lexicographer or made a clear and unambiguous disavowal of scope. *See Thorner*, 669 F.3d at 1365. Neither of those

exceptions apply here. For example, the disputed term is disclosed in the specification as one of ordinary skill in the art would expect and understand to be the plain and ordinary meaning of the term. *See Ex. H at 10:37-42; see also Madisetti Decl. at ¶¶ 34-36.*

VMware's construction improperly reads in limitations inconsistent with the description of the disputed term in the specification. Specifically, VMware's limitation "between two or more devices" is contrary to specification's broader disclosure. *See Ex. H at 10:37-38.*

Accordingly, IV's proposed construction of plain and ordinary meaning should be adopted.

C. "storing . . ." / "storing . . ." / "receiving . . ." / "determining . . ." / "determining . . ." / "sending . . ." / "using . . ." (claims 1 and 3)

VMware characterizes these claim elements as "ambiguous" and having a "significant antecedent basis problem that makes the claims impossible to decipher" in view of the constituent terms "physical interface(s)" and "tunnel identifier(s)" recited in claims 1 and 3. VMware offers no further clarification to support its assertion nor any statements explaining how or why specific instances of the constituent terms lack antecedent basis either individually or in combination.

There are no antecedent basis issues in claims 1 and 3. The disputed claims each begin by reciting generalized (i.e., unqualified by an article such as "a", "an", or "the") plural forms of each constituent term: "physical interfaces" and "tunnel identifiers." Ex. H at claims 1 and 3. Next, the claims sequentially recite the constituent terms using an indefinite article ("a") followed by a definite article ("the"): "a physical interface" / "the physical interface" and "a tunnel identifier" / "the tunnel identifier." *Id.* Lastly, the claims sequentially recite the constituent terms using an indefinite article followed by a qualified definite article that refers directly back to the preceding indefinite article: "determining . . . a physical interface and tunnel identifier" / "the determined [physical interface/tunnel identifier]" (claim 1), and "using . . . to identify a physical interface and tunnel identifier" / "the identified [physical interface/tunnel identifier]." *Id.*

As a matter of grammar, articles identify particular known (definite form) or unknown (indefinite form) instances of the terms they qualify. It is well-established under antecedent basis law that claim terms qualified by a definite article are interpreted to refer back to preceding instances of the term qualified by an indefinite article. *Baldwin Graphic Sys., Inc. v. Siebert, Inc.*, 512 F.3d 1338, 1342-43 (Fed. Cir. 2008) (recognizing “an indefinite article ‘a’ or ‘an’ in patent parlance” and “the subsequent use of definite articles ‘the’ or ‘said’ in a claim to refer back to the same claim term”). Further, reciting a term qualified by an indefinite article after a previous recitation of the term in generalized plural form does not necessarily implicate the singular term as an instance of the plural term, particularly when additional qualifying language is present. See *Medtronic, Inc. v. Guidant Corp.*, 465 F.3d 1360, 1377 (Fed. Cir. 2006) (finding no requirement that “a control signal” necessarily count as instance of previously recited “cardiac signals” and noting distinct qualifying language “control” and “cardiac”). The literal claim language as written is therefore entirely consistent with these established principles of antecedent basis law.

In addition to antecedent basis being satisfied on its face, the surrounding context of the claim language expressly clarifies which instances of “physical interfaces” and “tunnel identifiers” are associated with “customer lookup table/information” and the which instances are associated with “customer forwarding table(s)/information.” The claims when read in their entirety, in conjunction with antecedent basis being satisfied on its face, suffice to describe to a POSITA what the inventors were claiming. See *ACTV, Inc. v. Walt Disney Co.*, 346 F.3d 1082, 1088 (Fed. Cir. 2003).

As explained in IV’s Opening Brief, plain and ordinary meaning is the correct construction because the ’051 patent uses well-understood claim terms combined in a clear and straightforward manner. *IV Br.* at 45-46. At the time of the inventions of the ’051 patent, a POSITA would readily have understood the intended meaning of each disputed term. Thus, no further definition is necessary. See Madisetti Decl. at ¶¶ 37-39.

Despite there being no need to construe these omnibus claim elements beyond their plain and ordinary meaning, VMware characterizes its proposed constructions as “a simple way to resolve this problem that is fully supported by the specification.” This is not true and precisely why VMware’s proposed constructions are inappropriate i.e., because they violate the well-settled rule of claim construction that limitations described in particular embodiments in the specification should not be imported into a claim unless accompanied by statements of clear disavowal or otherwise required by the specification. *See Arlington Indus., Inc. v. Bridgeport Fittings, Inc.*, 632 F.3d 1246, 1254 (Fed. Cir. 2011).

VMware cites only examples in the specification where an incoming and an outgoing directionality is discussed. *VMware Br.* at 41-42. However, the specification confirms that the claim is broader than these examples. Indeed, the specification contemplates numerous instances where the claimed tunnel switching activity can operate in either or both directions. *See e.g.*, Ex. H at 12:6-20 (“in the opposite direction”); 11:51-54 (“tables 800 and 900 operate to switch transmissions in both directions”); 14:20-21 (“the tunnel functions as a bi-directional data pipe”); Figs. 6-7; 10B-11 (using bidirectional arrows and/or arrows in both directions to depict tunnel switching activity). Thus, VMware improperly asks the Court to read preferred embodiments out of the claim. *See Accent Packaging, Inc. v. Leggett & Platt, Inc.*, 707 F.3d 1318, 1326 (Fed. Cir. 2013) (“a claim interpretation that excludes a preferred embodiment from the scope of the claim is rarely, if ever, correct.”).

Other than its conclusory antecedent basis allegation concerning “physical interface(s)” and “tunnel identifier(s)” in particular, VMware offers no explanation as to why a POSITA would fail to understand any of the constituent claim terms recited in the elements it proposes to construe, nor why a POSITA would fail to understand how the combinations of said claim terms interrelate with each other. VMware has not met its burden of showing that the plain and ordinary meaning of the literal claim language is ambiguous or contradictory to the specification in any way.

D. “customer forwarding [table(s)/information]” (claims 1 and 3)

IV’s Proposed Construction	VMware’s Proposed Construction
“table(s) containing [a set/sets] of customer specific forwarding information” / “set(s) of customer specific forwarding information”	See construction of “storing a customer lookup table, the customer lookup table storing associations between physical interfaces and tunnel identifiers identifying tunnels for private networks and a plurality of customer forwarding tables” / “storing . . . customer forwarding information, . . . the customer forwarding information associating network addresses with physical interfaces and tunnel identifiers”

VMware does not specifically address IV’s proposed construction of the terms “customer forwarding table(s)” and “customer forwarding information” recited in claims 1 and 3.

Throughout much of the preliminary claim construction and subsequent meet-and-confer process, VMware’s original position was that it was relying on its separately proposed constructions for the terms discussed *supra* in Section 5(C) to construe the disputed terms at issue here (those separate, larger claim terms recite instances of the disputed terms within). However, VMware’s revised constructions no longer propose to alter or qualify the meaning of “customer forwarding table(s)” and “customer forwarding information” in any way. Rather, they focus on alleged antecedent basis issues associated with the “physical interface” and “tunnel identifier” terms. Although the parties did not specifically discuss dropping or agreeing upon a construction for this term, it appears VMware is suggesting by omission that it does not object to IV’s proposed constructions.

IV’s proposed constructions closely track the underlying claim language and seek only to add the word “set” or “sets” where appropriate to help one differentiate between singular versus plural instances of the disputed claim terms, and further to maintain consistency between claims 1 and 3. In addition to closely mirroring the literal claim language, IV’s constructions are consistent with a POSITA’s understanding of the terms in view of the specification. For example, the ’051 specification explains:

FIG. 8 is an embodiment of a customer lookup table **800**, and FIG. 9 is an embodiment of a set of customer forwarding tables **900**. Together, table **800** and set of tables **900** are suitable for switching a set of transmissions . . . Customer lookup table **800** is used as an

index into the correct customer forwarding table from the set of customer forwarding tables **900**.

Ex. H at 11:42-54; Figs. 8-9; *see also* Madisetti Decl. at ¶¶ 40-41.

The '818 Patent

The inventions of United States Patent No. RE 44,818 (“the ‘818 patent”) provide improved server virtualization solutions in demanding computing environments, such as data centers. Specifically, the ‘818 patent teaches the dynamic managing of networks, virtual machines, virtual storage, and virtual input/output resources using a more granular and flexible quality-of-service scheme that includes a hierarchical classification and enforcement mechanism for allocating bandwidth within those networks.

6. Disputed Terms in the '818 Patent

A. “hierarchical token bucket resource allocation” / “token(s)” (claims 1, 17, 30, 32 and 42)

IV’s Proposed Construction	VMware’s Proposed Construction
Plain and ordinary meaning	The specific class-based scheduling algorithm known in the art as the “hierarchical token bucket” / “token” as used in “hierarchical token bucket resource allocation”

Here, VMware seeks to cabin the well-understood concept of a hierarchical token bucket algorithm into a specific, extrinsic definition contrary to the intrinsic record. It is improper for VMware to substitute extrinsic evidence inconsistent with the broad disclosure found in the intrinsic record. *See Bell & Howell Document Mgmt. Prods. Co. v. Altek Sys.*, 132 F.3d 701, 706 (Fed. Cir. 1997) (“reliance on extrinsic evidence to interpret claims is proper only when the claim language remains genuinely ambiguous after consideration of the intrinsic evidence”).

As discussed in IV’s Opening Brief, the specification supports reading the term in accordance with its plain and ordinary meaning as understood in the context of the hierarchical QoS processes described therein. *See e.g.*, Ex. G at 2:1-18; 8:20-44; 9:61-10:67; 11:25-27; 11:51-52; 13:21-23; 13:40-44; Figs. 12-13. And, contrary to VMware’s proposed construction, the intrinsic record does not limit the term, let alone define it with a particular, extrinsic definition. Accordingly, IV’s proposed construction of plain and ordinary meaning is proper

since a POSITA would readily understand the intended meaning of the disputed term in view of the specification. *See* Madisetti Decl. ¶¶ 42-45.

As to VMware's limiting extrinsic construction, the inventors do not specifically cite or reference any extrinsic definition of a hierarchical token bucket scheme in the intrinsic record. Further, the claim drafters chose to expressly recite the term as using the indefinite article "a" ("a hierarchical token bucket resource allocation") as opposed to using definite phrasing such as "Hierarchical Token Bucket resource allocation" indicating that there was no intent to limit the definition to a specific, extrinsic approach or term of art. VMware asserts that "HTB is a "proper noun" that refers "specifically to the scheduling mechanism known as 'hierarchical token bucket'" and not anything else" and points to the following language from the '818 patent: "using scheduling and queuing methods such as hierarchical token bucket (HTB)." Ex. G at 9:64-65. However, VMware's argument appears to collapse the distinction between proper nouns and acronyms. Although acronyms are always capitalized, many acronyms do not stand for proper nouns and should not be capitalized in their written-out form. As is the case in the '818 patent, the specification first and foremost identifies "hierarchical token bucket" without capitalization as a common noun. The fact that the "(HTB)" acronym which follows uses capitalization is irrelevant because it is a general property of all acronyms.

Further, VMware's arguments contradict its own position. As it acknowledges in its Opening Brief, "HTB was invented or at least popularized by Martin Devera when he implemented it as a queuing discipline in the Linux operating system." VMware Br. at 42. VMware thus concludes that the term could only logically refer to "the very specific method of allocating bandwidth resources referred to in the art." *Id.* However, as evidenced by VMware's own extrinsic evidence citations, numerous implementations of a hierarchical token bucket algorithm have been introduced into the art since the initial usage of the concept as a Linux queuing mechanism. *Id.* (referencing Snoeren Decl. ¶¶ 107-112) (listing multiple hierarchical token bucket implementations including Linux implementation, WLAN implementation, container implementation, etc.). It is well-settled that the meaning of a disputed term (as

understood by a POSITA) is measured relative to the critical date of the invention at issue, not the inception date of the term in the prior art. *See Innova/Pure Water, Inc. v. Safari Water Filtration Sys., Inc.*, 381 F.3d 1111, 1116-17 (Fed. Cir. 2004) (“[a] court construing a patent claim seeks to accord a claim the meaning it would have to a person of ordinary skill in the art at the time of the invention.”).

Although the concept of a hierarchical token bucket algorithm may have originated and first been implemented as a specific Linux-based queuing mechanism, various additional implementations emerged over the years prior to the invention of the ’818 patent demonstrating that usage of the term in the art evolved and diversified over time following its initial inception. Even assuming, *arguendo*, that VMware’s cited reference was the way to implement a hierarchical token bucket when it was written, the patent’s broad disclosure demonstrates that it was not the only way at the time of invention. Accordingly, VMware’s characterization of the term as a “very specific method . . . referred to in the art” is misguided and aimed solely at narrowing the meaning of the term without specifically identifying any deficiencies in the claim language itself or associated passages in the specification. The appropriate construction of “hierarchical token bucket resource allocation” should be its plain and ordinary meaning as interpreted in light of the specification.

B. “enforcing . . .” / “receiv[e/ing] . . .” / “classify[ing] . . .” / “compar[e/ing] . . .” / “forward[ing] . . .” / “buffer[ing] . . .” (claims 1, 17, 30, 32, 33, 37-39, 42)

IV’s Proposed Construction	VMware’s Proposed Construction
Plain and ordinary meaning	“enforcing . . . across the physical [storage network] interface of the virtual I/O server” / “receiv[e/ing] in the virtual I/O server” / “classify[ing] in the virtual I/O server” / “compar[e/ing] in the virtual I/O server” / “forward[ing] in the virtual I/O server” / “buffer[ing] in the virtual I/O server”

Here, VMware asks the Court read in the limitation “virtual I/O server” into nearly every element of claims 1, 17, 30, 32, 33, 37-39, and 42 of the ’818 patent despite those words not appearing in the body of any asserted claim.

As support, VMware primarily relies on case law standing for the proposition that limitations that “repeatedly and consistently” describe the invention may limit claim scope. VMware Br. at 45. The facts here, however, are much different. In both *SkinMedica* and *Microsoft*, the patentee made statements describing two embodiments of a single term as necessarily distinct resulting in the court finding that the disputed term, in the context of the claim, excluded one of those embodiments.¹⁵

VMware’s proposed construction, however, does not seek to merely narrow the functional scope of an existing claim term based on its limiting characterization in the specification. Rather, it proposes to import wholesale new structural limitations into nearly every element of every claim. This violates the well-established precedent that examples from the specification should not be imported into the claim where the claim language is not limiting and/or is broader than embodiment(s) described in the specification. *See Arlington Indus., Inc.*, at 1254 (Fed. Cir. 2011).

This is particularly true when the importation of structural limitations is at issue, as is the case here. *See Paragon Sols., LLC v. Timex Corp.*, 566 F.3d 1075, 1084-85 (Fed. Cir. 2009) (rejecting importation of structural limitation into claims requiring “electronic positioning device” and “physiological monitor” to include separate physical structures; district court erroneously relied on examples in spec that separated structures as limiting).

Further, with respect to method claims 1, 32, and 42, they each recite “A method of facilitating management of input/output subsystems in a virtual input/output server, comprising” in the preamble. It is well-settled, however, that limitations from the preamble should not be imported into the body of the claim if the claim body defines a structurally complete invention. *See Georgetown Rail Equip. Co. v. Holland L.P.*, 867 F.3d 1229, 1236 (Fed. Cir. 2017) (preamble language not limiting if claim and spec otherwise recite structurally complete invention).

¹⁵ *C.f. SkinMedica, Inv. v. Histogen, Inc.*, 727 F.3d 1187, 1203-04 (Fed. Cir. 2013); *Microsoft Corp. v. Multi-TechSystems, Inc.*, 357 F.3d 1340, 1347-48 (Fed. Cir. 2004).

Furthermore, it is entirely proper to claim method body elements as a series of functional steps. *See Cox Commc 'ns, Inc. v. Sprint Commc 'n Co. LP*, 838 F.3d 1224, 1232 (Fed. Cir. 2016) ("the asserted claims are method claims, so it makes sense to define the inventive method as a series of functions").

VMware does not allege that any of the leading terms "enforce[e/ing]," "receiv[e/ing]," "classify[ing]," "compar[e/ing]," "forward[ing]," and "buffer[ing]" are in dispute, nor does it specifically identify any other literal claim terms as allegedly unclear or expressly redefined by the specification. This is precisely because the plain and ordinary meaning of the disputed terms is clear. All VMware seeks to do is narrow the claims by importing unclaimed structure from embodiments in the specification. This does nothing to clarify the meaning of the literal claim terms at issue other than imposing extraneous, unrecited limitations in nearly every claim element. *See, e.g., O2 Micro*, 521 F.3d at 1362 (district courts not required to construe every limitation present in asserted claims and should focus on claim terms for which the parties have presented a fundamental dispute).

The terms at issue are used clearly and unambiguously throughout the claims in a manner that is consistent with the specification. Thus, a POSITA would properly understand these terms and accord them their plain and ordinary meaning. *See* Madisetti Decl. at ¶¶ 46-47.

C. "maintain[ing] a connection over a network fabric" (claims 1, 17, 30, 32, 42)

IV's Proposed Construction	VMware's Proposed Construction
Plain and ordinary meaning	"maintaining a connection between the physical interface of the application server and the physical interface of the virtual I/O server over a network fabric"

VMware argues that two additional limitations should be read into the disputed term including "the physical interface of the application server" and "the physical interface of the virtual I/O server."

In support of its desire to read in the limitation "and the physical interface of the virtual I/O server," VMware first references the opinion of its expert Dr. Snoeren in stating that "a POSITA would understand that a physical interface connects to a network fabric" and further

concluding that a physical interface limitation must be read into the claim. It is well-settled that limitations from the specification should not be read into a claim unless the patentee acts as his or her own lexicographer or the specification otherwise makes clear statements of disavowal.

See Thorner, 669 F.3d at 1365; *Arlington*, 632 F.3d at 1254. Dr. Snoeren's own statement contradicts VMware's argument, since he relies directly on language from the '818 patent specification stating that “[i]n one implementation, the I/O fabric driver stack and I/O fabric PHY interface can employ the Reliable Connections (RC) supported by the Infiniband standard.” Ex. G at 3:9-11. As expressed literally by the patentees, the statement at issue merely describes one implementation and was clearly not intended to be limiting. Further, the statement appears to be describing characteristics of an exemplary I/O server as opposed to a network fabric itself.

Next, to support reading in the additional limitation “between the physical interface of the application server,” VMware argues that “Figure 2 depicts the protocol stack of an application server” and again references the opinion of its expert Dr. Snoeren in concluding that “the physical interface of the virtual I/O device connects to the network fabric.” VMware Br. at 47. Once again, VMware mischaracterizes the specification to suggest that Fig. 2 discloses “the protocol stack of an application server” where in fact the '051 specification states “(“virtual network interface, in one implementation, emulates an Ethernet NIC [and] plugs in at the bottom of the network stack.”) Ex. G at 4:9-13; *see also* at 4:2-4 (stating “[i]nserted into the network and storage protocol stacks are virtual interface drivers configured to intercept storage and network I/O messages” thus further contemplating the existence of multiple implementations of a protocol stack).

Also, despite the disputed term's inclusion of “maintain[ing] a connection over a network fabric,” VMware's argument focuses on the “network fabric” portion of the term in isolation. VMware further overlooks the presence and significance of the claim language that directly follows and qualifies the disputed term: “. . . to a virtual [storage/network] interface layer of an application server.” Taking the entire claim element in context, the patentees have already defined structural relationships between “a connection,” “a network fabric,” and “a virtual

interface layer of an application server.” Thus, it is unclear why it would be necessary or permissible to impose additional “physical interface” limitations when the claim drafters have already defined structure relating “a connection” to “a virtual interface.” It would have readily been understood by those having ordinary skill in the art in view of the specification that the term “connection” generally is not limited to a “physical connection” since connections can just as readily be logical or virtual, for example. Ex. G at 4:9-13 (“virtual network interface, in one implementation, emulates an Ethernet NIC [and] plugs in at the bottom of the network stack”); 6:29-32 (“[v]irtualization software in the virtual machine server abstracts the underlying hardware by creating an interface to virtual machines, which represent virtualized resources”); *see also* Madisetti Decl. at ¶¶ 48-49.

D. “virtual [network/storage network] interface layer of an application server” (claims 1, 17, 32, 42)

IV’s Proposed Construction	VMware’s Proposed Construction
Plain and ordinary meaning	“a virtual storage network interface to higher layers of the virtual node in an application server”/ “virtual network interface layer to higher layers of the virtual node in an application server” / “virtual interface layer to higher layers of the virtual node in an application server”

VMware’s proposed construction introduces two additional limitations to these terms: (1) “. . . to higher layers;” and (2) “. . . of the virtual node.” In doing so, VMware asks the Court to ignore two claim construction canons by asking the Court to read in preferred embodiments and to ignore the patentee’s use of the limitation in some claims but not others.

The dispute concerning “. . . to higher layers” narrows to VMware’s proposal to impose a plural “to higher layers” requirement in an attempt to read in limitations from the specification. Ex. G at 4:29-32 (“FIG. 2 illustrates the protocol stack and modules of an application server 102 according to one possible implementation of the invention.”) This disregards the well-established principle that a claim term is not limited to specific or even preferred embodiments of the specification absent a clear disclaimer or disavowal of claim scope. Worse, as described in IV’s Opening Brief, VMware’s construction would read out preferred embodiments. Ex. G at

4:14-32; 4:51-62; Fig. 2; *see also* IV Br. at 30. Such a construction is “rarely, if ever, correct.” *See Accent Packaging*, 707 F.3d at 1326.

With respect to VMware’s second limitation, “virtual node,” that term is recited in certain claims but not others. For instance, it appears in claims 32 and 42 but is absent from claims 1 and 17. Thus, the applicant knew how to claim “virtual node” when desired and how not to claim it when the intent was to not use the term. Underscoring this point, in some instances the virtual node language even directly qualifies the disputed claim term, such as in claim 42 reciting “wherein the virtual storage network interface layer is associated with a virtual storage node identifier.” In light of such explicit usage elsewhere in the claims, the patentees established a clear familiarity with the concept of a virtual node and further displayed the ability to recite said limitation expressly in the claims when intended. VMware’s proposed inclusion of the words “of the virtual node” as qualifying the virtual interface layer(s) is, therefore, another attempt to read additional limitations into a claim term contrary to the surrounding context. *See Hockerson-Halberstadt*, 183 F.3d at 1374 (“the context of the surrounding words of the claim also must be considered in determining the ordinary and customary meaning of those claims.”).

VMware further fails to even suggest that either constituent phrase: “virtual interface layer . . .” or “. . . of an application server” is itself ambiguous, or that the relationship between the two phrases is ambiguous, which further undermines their position. To the extent the ’818 specification does discuss the concept of nodes, there are examples in which application servers are referred to as instances of “nodes” generally, but not “virtual” nodes as VMware would suggest. *See e.g.*, Ex. G 5:37-39 (“For example, the monitor module 250 is operative to automatically discover nodes (e.g., other application servers 102, virtual I/O servers 106).” For at least this additional reason, VMware’s proposed inclusion of the language “of the virtual node” is inappropriate as it deviates from the literal words of the disputed claim language at issue and seeks only to read in limitations that are neither required nor suggested by the intrinsic record. Again, VMware does not even argue that the constituent phrases are vague either individually or in combination. *See Homeland Housewares*, 865 F.3d at 375 (“[t]he claim

construction inquiry . . . begins and ends in all cases with the actual words of the claim.”); *Renishaw*, 158 F.3d at 1248 (“[t]his is so because the claims define the scope of the right to exclude”).

E. Alleged Means Plus Function Elements (claim 17) (individually set forth in Ex. C)

VMware alleges that ten elements of claim 17 implicate § 112 ¶ 6 and are also invalid for failure to disclose sufficient structure. VMware Opening Br. at 49-50. In its brief, however, VMware dedicates less than a page to supporting these allegations, all of which are conclusory and unsupported by the intrinsic record. As discussed in detail in IV’s Opening Brief, Section C(5), these elements of claim 17 (listed in full in Ex. C) are not subject to § 112 ¶ 6 because: (1) “means for” is not present in any claims, (2) the prefatory language used does not consist of merely a nonce word followed by function, and (3) the claim language itself provides sufficient structure to avoid the application of § 112 ¶ 6. *See Fisher-Rosemount Sys.*, 2019 WL 6830806, at *15-*16.

VMware first claims that ‘module’ is a nonce word and the preceding qualifier ‘input/output virtualization’ term “fails to impart any structural significance.” This argument, however, overlooks the significance of the term ‘input/output virtualization module’ as a whole, the language of the claim, and the knowledge of one of skill in the art. For instance, an ‘input/output virtualization module’ connotes a distinct structure to one of skill in the art. *See Madisetti Decl.* at ¶ 50. The fact that in some cases when a module is claimed alone without any qualifiers it can be considered a nonce word is not applicable here because in claim 17 the term is qualified as an ‘input/output virtualization module.’ As the court in *Zeroclick* noted, it is inappropriate to apply § 112 ¶ 6 to an alleged nonce term when it is qualified with terms denoting structure to one of skill in the art. *See Zeroclick, LLC v. Apple Inc.*, 891 F.3d 1003, 1007-09 (Fed. Cir. 2018). While the claim at issue in *Zeroclick* was “user interface code,” the principal applies to the term ‘input/output virtualization module’ as well, because ‘input/output

‘virtualization’ is understood as a structural element to one of skill in the art, particularly in light of the claim read as a whole. *See* Madisetti Decl. at ¶ 50.

VMware’s argument that the claim elements lack structural support is further belied by the claim language preceding the ‘input/output virtualization module’ elements. The claim language itself includes “an input/output fabric interface,” “a storage network interface,” “one or more processors,” and “a memory.” *See* Ex. G at claim 17. Thus the ‘input/output virtualization module’ is tied—without leaving the four corners of the claim—to the above noted structure as understood by one of ordinary skill. *See* Madisetti Decl. at ¶ 50. This also indicates that application of § 112 ¶ 6 is not appropriate. *See Telcordia Techs., Inc. v. Cisco Sys., Inc.*, 612 F.3d 1365 (Fed. Cir. 2010).

Finally, as with the alleged means-plus-function terms of the ’686 patent, VMware here misapplies the standard for whether corresponding structure is disclosed. Assuming that the terms are found to be subject to § 112 ¶ 6—which they should not be—the specification does in fact disclose corresponding structure for these disputed terms, which IV noted in its Opening Brief at pg. 36-37, and its disclosure of proposed constructions (*see* Ex. C). For instance, the function claimed in element 1 of Ex. C is tied to structure in the specification including a network fabric, and a virtual storage network interface layer. The same holds true for elements 2-3 which are clearly linked to the structure of a physical network storage interface. *See* Ex. G at 3:4-7; Fig. 7. The remainder of the disputed terms (elements 4-10) either individually recite sufficient structure or refer back to the structural elements discussed above. *See, e.g.*, Ex. C (defining structure for disputed terms). This fact would be easily understood by one of skill in the art, which is the standard for the sufficiency of structure in a § 112 ¶ 6 analysis. *See Telcordia Techs.*, 612 F.3d at 1377 (“claim definiteness depends on the skill level of an ordinary artisan. Therefore, the specification need only disclose adequate defining structure to render the bounds of the claim understandable to an ordinary artisan.”); *see also, Intelligent Automation Design*, 2020 WL 486830, at *3-*5.

Conclusion

For the reasons stated herein and in IV's Opening Claiming Construction Brief, IV's proposed constructions should be adopted.

Dated: March 27, 2020

/s/ Robert R. Gilman
DEREK GILLILAND
State Bar No. 24007239
SOREY, GILLILAND & HULL, LLP
109 W. Tyler Street
Longview, Texas 75601
903.212.2822 (telephone)
903.212.2864 (facsimile)
derek@soreylaw.com

KARL RUPP
State Bar No. 24035243
NIX PATTERSON L.L.P.
1845 Woodall Rodgers Fwy. Suite 1050
Dallas, Texas 75201
972.831.1188 (telephone)
972.444.0716 (facsimile)
krupp@nixlaw.com

OF COUNSEL:

Paul J. Hayes
phayes@princelobel.com
Matthew Vella
mvella@princelobel.com
Robert R. Gilman
rgilman@princelobel.com
Jonathan DeBlois
jdeblois@princelobel.com
Alex Breger
abreger@princelobel.com
PRINCE LOBEL TYE LLP
One International Place, Suite 3700
Boston, MA 02110
Tel: (617) 456-8000
Fax: (617) 456-8100

COUNSEL for PLAINTIFFS

CERTIFICATE OF SERVICE

I hereby certify that on the 27th day of March, 2020, I electronically filed the foregoing with the Clerk of using the CM/ECF system which will send notification of such filing to all counsel of record.

/s/ Robert R. Gilman